

# Concept clustering and knowledge integration from a children's dictionary

Caroline Barrière and Fred Popowich

School of Computing Science, Simon Fraser University  
 Burnaby, BC, Canada, V5A 1S6  
 barriere,popowich@cs.sfu.ca

## Abstract

Knowledge structures called Concept Clustering Knowledge Graphs (CCKGs) are introduced along with a process for their construction from a machine readable dictionary. CCKGs contain multiple concepts interrelated through multiple semantic relations together forming a semantic cluster represented by a conceptual graph. The knowledge acquisition is performed on a children's first dictionary. The concepts involved are general and typical of a daily life conversation. A collection of conceptual clusters together can form the basis of a lexical knowledge base, where each CCKG contains a limited number of highly connected words giving useful information about a particular domain or situation.

## 1 Introduction

When constructing a Lexical Knowledge Base (LKB) useful for Natural Language Processing, the source of information from which knowledge is acquired and the structuring of this information within the LKB are two key issues. Machine Readable Dictionaries (MRDs) are a good source of lexical information and have been shown to be applicable to the task of LKB construction (Dolan et al., 1993; Calzolari, 1992; Copestake, 1990; Wilks et al., 1989; Byrd et al., 1987). Often though, a localist approach is adopted whereby the words are kept in alphabetical order with some representation of their definitions in the form of a template or feature structure. Efforts in finding connections between words is seen in work on automatic extraction of semantic relations from MRDs (Ahlsweide and Evens, 1988; Alshawi, 1989; Montemagni and Vanderwende, 1992). Additionally, efforts in finding words that are close seman-

tically is seen by the current interest in statistical techniques for word clustering, looking at co-occurrences of words in text corpora or dictionaries (Church and Hanks, 1989; Wilks et al., 1989; ?; Pereira et al., 1995).

Inspired by research in the areas of semantic relations, semantic distance, concept clustering, and using Conceptual Graphs (Sowa, 1984) as our knowledge representation, we introduce Concept Clustering Knowledge Graphs (CCKGs). Each CCKG will start as a Conceptual Graph representation of a trigger word and will expand following a search algorithm to incorporate related words and form a Concept Cluster. The concept cluster in itself is interesting for tasks such as word disambiguation, but the CCKG will give more to that cluster. It will give the relations between the words, making the graph in some aspects similar to a script (Schank and Abelson, 1975). However, a CCKG is generated automatically and does not rely on primitive s but on an unlimited number of concepts, showing objects, persons, and actions interacting with each other. This interaction will be set within a particular domain, and the trigger word should be a key word of the domain to represent. If that process would be done for the whole dictionary, we would obtain an LKB divided into multiple clusters of words, each represented by a CCKG. Then during text processing for example, a portion of text could be analyzed using the appropriate CCKG to find implicit relations and help understanding the text.

Our source of knowledge is the American Heritage First Dictionary<sup>1</sup> which contains 1800 entries and is designed for children of age six to eight. It is made for young people learning the structure and the basic vocabulary of their language. In comparison, an adult's dictionary is more of a reference tool which assumes knowl-

<sup>1</sup>Copyright ©1994 by Houghton Mifflin Company. Reproduced by permission from THE AMERICAN HERITAGE FIRST DICTIONARY.

edge of a large basic vocabulary, while a learner's dictionary assumes a limited vocabulary but still some very sophisticated concepts. Using a children's dictionary allows us to restrict our vocabulary, but still work on general knowledge about day to day concepts and actions.

In the following sections, we first present the transformation steps from the definitions into conceptual graphs, then we elaborate on the integration process, and finally, we close with a discussion.

## 2 Transforming Definitions

Our definitions may contain up to three general types of information, as shown in the examples in Figure 1.

- **description:** This contains genus/differentia information. Such information is frequently used for noun taxonomy construction (Byrd et al., 1987; Klavans et al., 1990; Barrière and Popowich, 1996).
- **general knowledge or usage:** This gives information useful in daily life, like how to use an object, what it is made of, what it looks like, etc.
- **specific example:** This presents a typical situation using the word defined and it involves specific persons and actions.

**Cereal** is a kind of food. [description]  
 Many **cereals** are made from corn, wheat, or rice. [usage]  
 Most people eat **cereal** with milk in a bowl. [usage]

**Ash** is what is left after something burns. [usage]  
 It is a soft gray powder. [description]  
 Ray watched his father clean the **ashes** out of the fireplace. [example]

Figure 1: Example of definitions

The information given by the description and general knowledge will be used to perform the knowledge integration proposed in section 3. The specific examples are excluded as they tend to involve specific concepts not always deeply related to the word defined.

Our processing of the definitions results in the construction of a special type of conceptual graph which we call a *temporary graph*. The set of relations used in temporary graphs come from three sources. Table 1 shows some examples for each type.

1. the set of closed class words, ex: *of, to, in, and*;
2. relations extracted via defining formulas ex: *part-of, made-of, instrument*; defining formulas correspond to phrasal patterns that occur often through the dictionary suggesting particular semantic relations (ex. A is a part of B) (?; Dolan et al., 1993).

3. the relations that are extracted from the syntactic structure of a sentence, ex: *subject, object, goal, attribute, modifier*.

As some relations are defined using the closed class words, and many of those words are ambiguous, the resulting graph will itself be ambiguous. This is the main reason for calling our graphs *temporary* as we assume a conceptual graph, the ultimate goal of our translation process, should contain a restricted set of well-defined and non-ambiguous semantic relations. For example, *by* can be a relation of manner (*by chewing*), time (*by noon*) or place (*by the door*). By keeping the preposition itself within the temporary graph, we delay the ambiguity resolution process until we have gathered more information and we even hopefully avoid the decision process as the ambiguity might later be resolved by the integration process itself.

	closed class words	temporary graph
np:np[A],prep[B],np[C]		[A]->(B)->[C]
apple on the table		[apple]->(on)->[table]
defining formulas		temporary graph
A is used to B		[B]->(instrument)->[A]
A is a part of B		[A]->(part-of)->[B]
A is a place where B		[B]->(loc)->[A]
syntactic pattern		temporary graph
s:np[A],vp[B]		[B]->(agent)->[A]
John eats		[eat]->(agent)->[John]
	vp:vp[A],inf_vp[B]	[A]->(goal)->[B]
	eat to grow	[eat]->(goal)->[grow]

Table 1: Examples of relations found in sentences and their corresponding temporary graphs

## 3 Knowledge Integration

This section describes how given a trigger word, we perform a series of forward and backward searches in the dictionary to build a CCKG containing useful information pertaining to the trigger word and to closely related words. The primary building blocks for the CCKG are the temporary graphs built from the dictionary definitions of those words using our transformation process described in the previous section. Those temporary graphs express similar or related ideas in different ways and with different levels of detail. As we will try to put all this information together into one large graph, we must first find what information the various temporary graphs have in common and then join them around this common knowledge.

To help us build this CCKG and perform our integration process, we assume two main knowledge structures are available, a concept hierarchy

and a relation hierarchy, and we assume the existence of some graph operations. The concept hierarchy concentrates on nouns and verbs as they account for three quarters of the dictionary definitions. It has been constructed automatically according to the techniques described in (Barrière and Popowich, 1996). The relation hierarchy was constructed manually. A rich hierarchical structure between the set of relations is essential to the graph matching operations we use for the integration phase.

As we are using the conceptual graph formalism to represent our definitions, we can use the graph matching operations defined in (Sowa, 1984). The two operations we will need are the Maximal Common Subgraph algorithm and the Maximal Join algorithm.

### 3.1 Maximal Common Subgraph

The maximal common subgraph between two graphs consists of finding a subgraph of the first graph that is isomorphic to a subgraph of the second graph. In our case, we cannot often expect to find two graphs that contain an identical subgraph with the exact same relations and concepts. Ideas can be expressed in many ways and we therefore need a more relaxed matching schema. We describe a few elements of this “relaxation” process and illustrate them by an example in Figure 2.

<p>(1) John makes a nice drawing on a piece of paper with the pen.</p> <pre>[make]-&gt;(sub)-&gt;[John]     -&gt;(obj)-&gt;[drawing]-&gt;(att)-&gt;[nice]     -&gt;(on)-&gt;[piece]-&gt;(of)-&gt;[paper]     -&gt;(with)-&gt;[pen]</pre>																	
<p>(2) John uses the big crayon to draw rapidly on the paper.</p> <pre>[draw]-&gt;(sub)-&gt;[John]     -&gt;(on)-&gt;[paper]     -&gt;(instrument)-&gt;[crayon]     -&gt;(manner)-&gt;[rapidly]</pre>																	
<p>MAXIMAL COMMON SUBGRAPH:</p> <pre>[make(draw)]-&gt;(sub)-&gt;[John]     -&gt;(obj)-&gt;[drawing]     -&gt;(on)-&gt;[piece]-&gt;(of)-&gt;[paper]     -&gt;(instrument)-&gt;[label-1]</pre>																	
<table border="1"> <thead> <tr> <th>Relaxation method</th> <th>graph1</th> <th>graph2</th> </tr> </thead> <tbody> <tr> <td>Semantic distance</td> <td>pen</td> <td>crayon</td> </tr> <tr> <td>Relation subsumption</td> <td>with</td> <td>instrument</td> </tr> <tr> <td>Predictable meaning shift</td> <td>drawing</td> <td>draw</td> </tr> <tr> <td>Relation transitivity</td> <td>piece of paper</td> <td>paper</td> </tr> </tbody> </table>			Relaxation method	graph1	graph2	Semantic distance	pen	crayon	Relation subsumption	with	instrument	Predictable meaning shift	drawing	draw	Relation transitivity	piece of paper	paper
Relaxation method	graph1	graph2															
Semantic distance	pen	crayon															
Relation subsumption	with	instrument															
Predictable meaning shift	drawing	draw															
Relation transitivity	piece of paper	paper															
<p>MAXIMAL JOIN:</p> <pre>[make(draw)]-&gt;(sub)-&gt;[John]     -&gt;(obj)-&gt;[drawing]-&gt;(att)-&gt;[nice]     -&gt;(on)-&gt;[piece]-&gt;(of)-&gt;[paper]     -&gt;(instrument)-&gt;[label-1]     -&gt;(manner)-&gt;[rapidly]</pre>																	

Figure 2: Example of “relaxed” maximal common subgraph and maximal join algorithms

**Semantic distance between concepts.** In the maximal common subgraph algorithm proposed by (Sowa, 1984), two concepts (C1,C2) could be matched if one subsumed the other in the concept hierarchy. We can relax that criteria to match two concepts when a third concept C which subsumes C1 and C2 has a high enough degree of informativeness (Resnik, 1995). The concept hierarchy can be useful in many cases, but it is generated from the dictionary and might not be complete enough to find all similar concepts.

In the example of Figure 2, when using the concept hierarchy to establish the similarity between *pen* and *crayon*, we find that one is a subclass of *tool* and the other of *wax*, both then are subsumed by the general concept *something*. We have reached the root of the noun tree in the concept hierarchy and this would give a similarity of 0 based on the informativeness notion.

We extend the subsumption notion to the graphs. Instead of finding a concept that subsumes two concepts, we will try finding a common subgraph that subsumes the graph representation of both concepts. In our example, *pen* and *crayon* have a common subgraph `[write]->(inst)->[]`. The notion of semantic distance can be seen as the informativeness of the subsuming graph. The resulting maximal common subgraph as shown in Figure 2 contains the concept *label-1*. This label is associated to a covert category as presented in (Barrière and Popowich, 1996). We can update the concept hierarchy and add this *label-1* as a subclass of *something* and a superclass of *pen* and *crayon*. It expresses a concept of “writing instrument”.

**Relation subsumption.** Since we have a relation hierarchy in addition to our concept hierarchy, we can similarly use subsumption to match two relations. In Figure 2, *with* is subsumed by *instrument*, and by mapping them, we disambiguate *with* from corresponding to another semantic relation, such as *possession* or *accompaniment*. This is a case where an ambiguous preposition left in the temporary graph is resolved by the integration process.

**Predictable meaning shift.** A set of lexical implication rules were developed by (Ostler and Atkins, 1992) for relating word senses. Based on them, we are developing a set of graph matching rules. Figure 2 exemplifies one of them where two graphs containing the same word (or morphologically related), here *draw* and *drawing*, used as different parts of speech can be related.

**Relation transitivity.** Some relations, like *part-of*, *in*, *from* can be transitive. For example, we can map a graph that contains a concept A in a certain relation to concept B onto another graph where concept A is in the same relation with a *part* or a *piece* of B as exemplified in Figure 2. Transitivity in relations is in itself a challenging area of study (Cruse, 1986) and we have only begun to explore it.

### 3.2 Maximal Join

The basic operation for the integration of temporary graphs is the maximal join operation where a union of two graphs is formed around their maximal common subgraph using the most specific concepts of each. We just saw how to relax the maximal common subgraph operation and we will perform the join around that “relaxed” subgraph. Figure 2 shows the result of the maximal join. The join operation allows us to bring new concepts into a graph by finding relations with existing concepts, as well as bringing new relations between existing concepts.

### 3.3 Integration process

Given the concept hierarchy, relation hierarchy and graph matching operations, we now describe the two major steps required to integrate all the temporary graphs into a CCKG.

**TRIGGER PHASE.** Start with a central word, a keyword for the subject of interest that becomes the trigger word. The temporary graph built from the trigger word forms the initial CCKG. To expand its meaning, we want to look at the important concepts involved and use their respective temporary graphs to extend our initial graph. We deem words in the definition to be important if they have a large semantic weight.

The semantic weight of a word or its informativeness can be related to its frequency (Resnik, 1995). Here, we calculate the number of occurrence of each word within the definitions of nouns and verbs in our dictionary. The most frequent word “a” occurs 2600 times among a total of 38000 word occurrences. Only 1% of the words occur more than 130 times, 5% occur more than 30 times but over 60% occur less than 5 times.

Ordering the dictionary words in terms of decreasing number of occurrences, the top 10% of these words account for 75% of word occurrences. For our current investigation, we propose this as the division between semantically significant words, and semantically insignificant ones. So a word from the dictionary is deemed to be semantically significant if it occurs less than 17 times.

Note that constraining the number of semantically significant words is important in limiting the exploration process for constructing the concept cluster, as we shall soon see.

**Trigger forward:** Find the semantically significant words part of the CCKG, and join their respective temporary graph to the initial CCKG.

**Trigger backward:** Find all the words in the dictionary that use the trigger word in their definition and join their respective temporary graph to the CCKG.

Instead of a single trigger word, we now have a cluster of words that are related through the CCKG. Those words form the *concept cluster*.

**EXPANSION PHASE.** We try finding words in the dictionary containing many concepts identical to the ones already present in the CCKG but perhaps interacting through different relations allowing us to create additional links within the set of concepts present in the CCKG. Our goal is to create a more interconnected graph rather than sprouting from a particular concept. For this reason, we establish a graph matching threshold to decide whether we will join a new graph to the CCKG being built. We set this threshold empirically: the maximal common subgraph between the CCKG and the new temporary graph must contain at least three concepts connected through two relations.

**Expansion forward:** For each semantically significant word in the CCKG, not already part of the concept cluster, find the maximal common subgraph between its temporary graph and the CCKG. If matching surpasses the graph matching threshold, perform integration (maximal join operation) and add the word in the concept cluster. Continue forward until no changes are made.

**Expansion backward:** Find words in the dictionary whose definitions contain the semantically significant words from the concept cluster. For each possible new word, perform the maximal common subgraph between its temporary graph and the CCKG. Again, if matching is over the graph matching threshold, perform integration and add the word in the concept cluster. Continue until no changes are made.

We can set a limit to the number of steps in the expansion phase to ensure its termination. However in practice, after two or three steps forward

or backward, the maximal common subgraphs between the new graphs and CCKG do not exceed the graph matching threshold and thus are not added to the cluster, terminating the expansion.

### 3.4 Example of integration

Figure 3 shows the starting point of an integration process with the trigger word (TW) *letter*, its definition, its temporary graph (TG), the concept cluster (CC) containing only the trigger word, and the CCKG being the same as the temporary graph. Then we show the trigger forward phase. The number of occurrences (NOcc) of each word present in the definition of *letter* is given. Using the criteria described in the previous section, only the word *message* is a semantically significant word (SSW). We then see the definition of *message*, the new concept cluster and the resulting CCKG.

The trigger backward phase, would incorporate the temporary graphs for *address*, *mail*, *post office* and *stamp*. The expansion forward phase would further add the temporary graphs for the semantically significant words: *{send, package}* during the first step and then would terminate with the second step as no more semantically significant words not yet explored have a maximal common subgraph with the CCKG that exceeds the graph matching threshold. The expansion backward would finally add the TGs for *card* and *note*, again terminating after two steps.

The resulting cluster is: *{letter, message, address, mail, post office, stamp, send, package, card, note}*. The resulting CCKG shows the interaction between those concepts which summarizes general knowledge about how we use those concepts together in a daily conversation: we go to the post office to mail letters, or packages; we write letters, notes and card to send to people through the mail, etc. Having such clusters and such knowledge of the relationship between words as part of our lexical knowledge base can be useful to understand or even generate a text containing the concepts involved in the cluster.

## 4 Discussion

Through this paper, we showed the multiple steps leading us to the building of Concept Clustering Knowledge Graphs (CCKGs). Those knowledge structures are built within the Lexical Knowledge Base (LKB), integrating multiple parts of the LKB around a particular concept to form a cluster and express the multiple relations among the words in that cluster. The CCKGs could be either permanent or temporary structures depending on

### STARTING POINT:

**TW:** *letter*  
**Def:** A *letter* is a message you write on paper.  
**TG:** same as CCKG  
**CC:** *{letter}*  
**CCKG:** [write]->(obj)->[message(letter)]  
  ->(subj)->[person:you]  
  ->(on)->[paper]

### TRIGGER FORWARD:

**NOccs:** you:280, paper:42, write:31,  
  message:7  
**SSWs:** message  
**Def:** A *message* is a group of words that is sent  
  from one person to another.  
  Many people send **messages** through the mail.  
**CC:** *{letter, message}*  
**CCKG:**  
  [word:group(message(letter))]  
  <-(obj)<-[write]->(sub)->[person:you]  
  ->(on)->[paper]  
  <-(obj)<-[send]->(subj)->[person:many]  
  ->(from)->[person:one]  
  ->(to)->[person:another]  
  ->(through)->[mail]

Figure 3: Trigger forward from *letter*.

the application using the LKB. For example, for a text understanding task, we can build before hand the CCKGs corresponding to one or multiple keywords from the text. Once built, the CCKGs will help us in our comprehension and disambiguation of the text.

By using the American Heritage First Dictionary as our source of lexical information, we were able to restrict our vocabulary to result in a project of reasonable size, dealing with general knowledge about day to day concepts and actions. The ideas explored using this dictionary can be extended to other dictionaries as well, but the task might become more complex as the definitions in adult's dictionaries are not as clear and usage oriented. In fact, an LKB built from a children's dictionary could be seen as a starting point from which we could extend our acquisition of knowledge using text corpora or other dictionaries. Certainly, if we envisage applications trying to understand children's stories or help in child education, a corpora of texts for children would be a good source of information to extend our LKB.

The graph operations (maximal common subgraph and maximal join) defined on conceptual graphs, and adapted here, play an important role in our integration process toward a final CCKG. Graph matching was also suggested as an alternative to taxonomic search when trying to establish semantic similarity between concepts. As well, by putting a threshold on the graph matching process, we were able to limit the expansion of our

clustering, as we can decide and justify the incorporation of a new concept into a particular cluster.

Many aspects of the concept clustering and knowledge integration processes have already been implemented and it will soon be possible to test the techniques on different trigger words using different thresholds to see how they effect the quality of the clusters.

Clustering is often seen as a statistical operation that puts together words “somehow” related. Here, we give a meaning to their clustering, we find and show the connections between concepts, and by doing so, we build more than a cluster of words. We build a knowledge graph where the concepts interact with each other giving important implicit information that will be useful for Natural Language Processing tasks.

## 5 Acknowledgments

The authors would like to thank the anonymous referees for their comments and suggestions. This research was supported by the Institute for Robotics and Intelligent Systems. The authors also want to thank Petr Kubon for his many comments on the paper.

## References

[Ahlsweide and Evens1988] T. Ahlsweide and M. Evens. 1988. Generating a relational lexicon from a machine-readable dictionary. *International Journal of Lexicography*, 1(3):214–237.

[Alshawi1989] H. Alshawi. 1989. Analysing the dictionary definitions. In Bran Boguraev and Ted Briscoe, editors, *Computational Lexicography for Natural Language Processing*, chapter 7, pages 153–170. Longman Group UK Limited.

[Barrière and Popowich1996] Caroline Barrière and Frederick Popowich. 1996. Building a noun taxonomy from a children’s dictionary. In *Euralex’96*. To be presented at Euralex’96, Göteborg, Sweden, August 96.

[Brown et al.1992] P. Brown, V.J. Della Pietra, P.V. deSouza, J.C. Lai, and R.L. Mercer. 1992. Class-based n-gram models of natural language. *Computational Linguistics*, 18(4):467–480.

[Byrd et al.1987] R.J. Byrd, N. Calzolari, M. Chodorow, J. Klavans, M. Neff, and O. Rizk. 1987. Tools and methods for computational lexicology. *Computational Linguistics*, 13(3-4):219–240.

[Calzolari1992] N. Calzolari. 1992. Acquiring and representing semantic information in a lexical knowledge base. In J. Pustejovsky and S. Bergler, editors, *Lexical Semantics and Knowledge Representation : First SIGLEX Workshop*, chapter 16, pages 235–244. Springer-Verlag.

[Church and Hanks1989] K. Church and P. Hanks. 1989. Word association norms, mutual information and lexicography. In *Proceedings of the 27th Annual meeting of the Association for Computational Linguistics*, pages 76–83, Vancouver, BC.

[Copestake1990] A.A. Copestake. 1990. An approach to building the hierarchical element of a lexical knowledge base from a machine readable dictionary. In *Proceedings of the Workshop on Inheritance in Natural Language Processing, Tilburg*.

[Cruse1986] D.A. Cruse. 1986. *Lexical Semantics*. Cambridge University Press.

[Dolan et al.1993] W. Dolan, L. Vanderwende, and S. D. Richardson. 1993. Automatically deriving structured knowledge bases from on-line dictionaries. In *The First Conference of the Pacific Association for Computational Linguistics*, pages 5–14, Harbour Center, Campus of SFU, Vancouver, April.

[Klavans et al.1990] Judith Klavans, M. S. Chodorow, and N. Wacholder. 1990. From dictionary to knowledge base via taxonomy. In *Proceedings of the 6th Annual Conference of the UW Centre for the New OED: Electronic Text Research*, pages 110–132.

[Montemagni and Vanderwende1992] S. Montemagni and L. Vanderwende. 1992. Structural patterns vs. string patterns for extracting semantic information from dictionaries. In *Proc. of the 14<sup>th</sup> COLING*, pages 546–552, Nantes, France.

[Ostler and Atkins1992] N. Ostler and B.T.S. Atkins. 1992. Predictable meaning shift: Some linguistic properties of lexical implication rules. In J. Pustejovsky and S. Bergler, editors, *Lexical Semantics and Knowledge Representation : First SIGLEX Workshop*, chapter 7, pages 87–100. Springer-Verlag.

[Pereira et al.1995] F. Pereira, N. Tishby, and L. Lee. 1995. Distributional clustering of english words. In *Proc. of the 33<sup>th</sup> ACL*, Cambridge, MA.

[Resnik1995] Philip Resnik. 1995. Using information content to evaluate semantic similarity in a taxonomy. In *Proc. of the 14<sup>th</sup> IJCAI*, volume 1, pages 448–453, Montreal, Canada.

[Schank and Abelson1975] R. Schank and R. Abelson. 1975. Scripts, plans and knowledge. In *Advance papers 4th Intl. Joint Conf. Artificial Intelligence*.

[Sowa1984] J. Sowa. 1984. *Conceptual Structures in Mind and Machines*. Addison-Wesley.

[Wilks et al.1989] Y. Wilks, D. Fass, G-M Guo, J. McDonald, T. Plate, and B. Slator. 1989. A tractable machine dictionary as a resource for computational semantics. In Bran Boguraev and Ted Briscoe, editors, *Computational Lexicography for Natural Language Processing*, chapter 9, pages 193–231. Longman Group UK Limited.